Observing Bag Gain in JPEG Batch Steganography

Edgar Kaziakhmedov
Department of ECE
Binghamton University
ekaziak1@binghamton.edu

Eli Dworetzky
Department of ECE
Binghamton University
edworet1@binghamton.edu

Jessica Fridrich
Department of ECE
Binghamton University
fridrich@binghamton.edu

***Abstract*—The bag gain relates to a gain in security due to spreading payload among multiple covers when the steganographer maintains a positive communication rate. This gain is maximal for a certain optimal bag size, which depends on the embedding method, payload spreading strategy, communication rate, and the cover source. Originally discovered and analyzed in the spatial domain, in this paper we study this phenomenon for JPEG images across quality factors. Our experiments and theoretical analysis indicate that the bag gain is more pronounced for senders that assign payloads based on image content more aggressively and maintain a fixed payload per bag in terms of bits per DCT rather than per non-zero AC DCT. We also observe a larger bag gain for higher JPEG qualities.**

***Index Terms*—Batch steganography, pooled steganalysis, JPEG, bag gain**

I. Introduction

In batch steganography, the sender spreads the secret payload among multiple cover images (a bag) to decrease the chances of being detected by the Warden. The Warden pools evidence from the same bag of images to detect the use of steganography, a process known as pooled steganalysis. This is achieved by fusing soft outputs of Warden's Single-Image Detector (SID) applied to each image in the bag. Batch steganography and pooled steganalysis has been introduced by Ker in 2006 [1] and has since been a subject of intense research [2], [3], [4], [5], [6], [7], [8], [9], [10], [11], [12], [13], [14].

With the introduction of content-adaptive steganography researchers began studying payload spreading strategies that allocate chunks of the secret payload based on how detectable the embedding is in each image. The authors of [7] considered an omniscient Warden pooling optimally using the likelihood ratio test derived from a Gaussian model of Warden's detector output and three heuristic batch senders: the Image Merging Sender (IMS) and the Distortion / Detectability Limited Senders (DiLS / DeLS). The paradigm of Gaussian embedding with Gaussian pixel model has been extended to the IMS and its improved version called Adaptive Batch size Image Merging sender (AdaBIM) in [8], [13]. Detector-informed spreading strategies called Minimum Deflection Sender (MDS) and Shift Limited Sender (SLS) with an omniscient Warden were studied in the spatial domain in [12]. The authors also reported on a new phenomenon called the bag gain. Batch steganography in the JPEG domain for a Warden unaware of the spreading strategy was studied in [11].

WIFS'2023, December 4-7, 2023, Nuremberg, Germany. 978-1-7281-9930-6/20/$31.00

Figure 1 is an illustrative example of the bag gain showing the accuracy of Warden's pooled detector as a function of number of images sent (the bag size B). The pooled detector's accuracy initially decreases with increasing B, then levels off, and eventually increases as the square-root law [15] inevitably engages since the sender maintains a positive communication rate. The bag gain is the maximal drop in detectability. It has been observed in the spatial domain for all batch senders studied in [12], for many types of SIDs and pooled detectors, and for differently informed Wardens. The bag gain was analyzed and explained in [16] with a simple source model. In particular, it was shown to robustly manifest for batch senders that, loosely speaking, embed larger payloads in hard-to-steganalyze images and small payloads in easy-to-steganalyze images. Since the bag gain occurs for bag sizes that can be used in practice, the steganographer can decrease the chances of being caught simply by selecting a bag size that is neither too big nor too small.

In general, the manifestation of the bag gain depends on many factors, including the average communication rate, the response of the SID used for pooling, the aggressiveness of the batch sender to assign payloads, and the cover source. The work of Zakaria et al. [11], for example, does not show the bag gain because the results are reported over multiple batch senders and rates that were too large. To the best of knowledge of the authors, the bag gain has not been studied in the JPEG domain, which is the main focus of this paper.

In the next section, we describe batch steganography and pooled steganalysis in a more formal manner and introduce all necessary concepts and terminology. Section III describes the batch senders and Warden's poolers studied in this paper. The setup of our experiments and implementation details appear in Section IV with the results of all experiments, their interpretation, discussion, and analysis laid out in Section V. The paper is concluded in Section VI.

II. Batch steganography and pooled steganalysis

A cover bag of size B, $\mathbf{X} = (X_0^{(1)}, \dots, X_0^{(B)})$, is formed by independently selecting B cover images $X_0^{(1)}, \dots, X_0^{(B)}$ from a cover source of equally sized JPEG images, each with a total of N DCT coefficients. We assume that the steganographer maintains a fixed communication rate r in

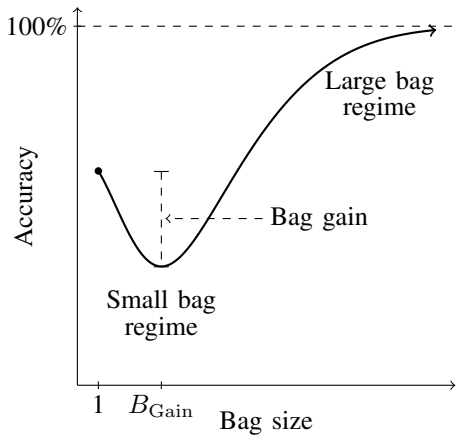

Figure 1. An illustrative example of the bag gain. Pooled detector's accuracy is shown as a function of bag size B. The steganographer can gain security by spreading payload among B_{Gain} cover images.

terms of either bits per DCT coefficient (bpc) or bits per non-zero AC DCT coefficient (bpnzac). This assumption is reasonable as a steganographic channel is likely to be used repetitively in practice. For a fixed positive rate r, the sender will eventually be caught due to the square root law.

In this paper, the relative payload α_i to be embedded in $X_0^{(i)}$ will always be measured in terms of bpc. A batch spreading strategy is a mapping that assigns payload chunks α_i bpc to all images in $\mathbf{X}$ so that the payload constraint is satisfied. Denoting the number of non-zero AC DCTs in $X_0^{(i)}$ as $N_i \leq N$, the payload constraint is

$$\sum_{i=1}^{B} \alpha_i = r \times c(\mathbf{X}), \tag{1}$$

where $c(\mathbf{X}) = B$ for rate r in bpc and $c(\mathbf{X}) = \sum_{i=1}^{B} N_i/N$ for r in bpnzac. The steganographer produces the ith stego image $X_{\alpha_i}^{(i)}$ by embedding $X_0^{(i)}$ with payload of length α_i.

Given an intercepted bag of B images $\mathbf{Y} = (Y^{(1)}, \ldots, Y^{(B)})$, the Warden infers whether steganography is being used by performing the following composite hypothesis test:

$$\begin{aligned} \mathcal{H}_0 &: \quad r = 0 \\ \mathcal{H}_1 &: \quad r > 0. \end{aligned} \tag{2}$$

The Warden "pools" the evidence $\mathbf{Y}$ together by using a pooled detector (or "pooler"). We assume the Warden's decision is solely informed by the collection of outputs of a SID, which is a mapping d that assigns to each image a scalar referred to as the soft output (or response) of the detector. Formally, the Warden's pooler is of the form $\pi : \mathbb{R}^B \to \mathbb{R}$, and she infers whether the sender uses steganography by computing $d(Y^{(i)})$ for all $i = 1, \ldots, B$ and comparing $\pi(d(Y^{(1)}), \ldots, d(Y^{(B)}))$ against a threshold determined by some application-dependent requirements.

III. Batch senders and poolers

We restrict our study to the IMS [7] and the MDS [12], [16]. The IMS is a well studied sender, which is a natural extension of steganography from a single image to a bag. It considers the entire bag as one large image and lets the embedding algorithm allocate the payload. The MDS was included because it is the most amenable to analysis within the context of a statistical model of the SID [16] and also the appendix. The model will help us understand, explain, and analyze trends observed in experiments in Section V.

A. MDS

The MDS assumes the Warden's hypothesis test has the following form (2):

$$\begin{aligned} \mathcal{H}_0 &: \quad d(Y^{(i)}) \sim \mathcal{N}(\mu_i, \sigma^2) \quad \text{for all } i \\ \mathcal{H}_1 &: \quad d(Y^{(i)}) \sim \mathcal{N}(\mu_i + s_i(\alpha_i), \sigma^2) \quad \text{for all } i, \end{aligned} \tag{3}$$

where $Y^{(i)}$ are the images from a bag under inspection by the Warden and α_i is the payload possibly residing in the ith image. The distribution of Warden's detector under $\mathcal{H}_0$ is over acquisitions of the same scene with the same camera and settings. The Gaussianity is heuristically justified by the independent heteroscedastic acquisition noise model [17] and the fact that d can be linearized on a small neighborhood of the noise-free scene. Furthermore, notice the additional simplifying assumption in (3) that the variance does not depend on i or α_i.

Assuming the parameters of the distributions in this hypothesis test are known to the Warden, the most powerful pooled detector is the likelihood ratio test. The detectability of steganography in a single bag is determined by the deflection coefficient

$$\Delta^2(\mathbf{X}) = \sum_{i=1}^{B} \frac{s_i^2(\alpha_i)}{\sigma^2} = \sum_{i=1}^{B} \frac{(\varrho_i(\alpha_i) - \varrho_i(0))^2}{\sigma^2}, \tag{4}$$

where

$$\varrho_i(\alpha) = \mathbb{E}[d(X_\alpha^{(i)}) | X_0^{(i)}] \tag{5}$$

is the so-called response curve of ith cover image for detector d, with the expectation taken over embedding $X_0^{(i)}$ with random messages of length α and stego keys. Note that $\varrho_i(\alpha)$ is assumed to be independent of the specific cover acquisition $X_0^{(i)}$.

Assuming that the sender uses the same SID d as the Warden,[1] the MDS minimizes the power of Warden's most powerful detector by selecting α_i that minimize $\Delta^2(\mathbf{X})$ subject to the payload constraint (1).

B. Poolers

In this paper, we use two types of SIDs: a binary classifier trained on a random uniform mixture of payloads (in terms of bpnzac) and a quantitative detector that returns an estimate of the length of the hidden payload.

We have investigated many different types of poolers but report only on their subset that performed the best. As the

[1] In reality and also in this paper, the detector used by the sender and the Warden do not match, hence the MDS will not be guaranteed to be optimal w.r.t. Warden's detector.

simplest uninformed pooler, we consider the average of soft SID outputs on analyzed images $Y^{(i)}$:

$$\pi_{\text{AVG}}(\mathbf{Y}) = \frac{1}{B}\sum_{i=1}^{B} d(Y^{(i)}). \tag{6}$$

For a Warden aware of the spreading strategy and rate r, we also use the correlator

$$\pi_{\text{COR}}(\mathbf{Y}) = \sum_{i=1}^{B} d(Y^{(i)})\hat{\alpha}_i, \tag{7}$$

where $\hat{\alpha}_i$ is the payload that might reside in ith image estimated by the Warden from the images at hand. For the IMS with J-UNIWARD [18], the payloads estimated from the stego images are nearly identical to the payloads computed from covers. This small estimation error has basically no effect on the detection accuracy. For the detector-aware MDS, the payloads need to be estimated from the response curves of the sender's detector. Since the Warden will generally not have access to this detector, there will always be an estimation error. Per our experiments, and in line with the findings reported in [12], this estimation error has a negligible effect on the detection accuracy. Hence, in all our experiments, we assume the worst case scenario for the sender and set $\hat{\alpha}_i = \alpha_i$.

We have also experimented with the max pooling strategy $\pi_{\text{MAX}}(\mathbf{Y}) = \max_i d(Y^{(i)})$ as well as machine learning poolers trained as Gaussian support vector machines for each bag size separately but do not report on them because the max pooler performed very poorly, while the learned poolers performed essentially the same as the average and the correlator when adding the payloads to the feature vector formed by outputs of the SID.

The implementation details of the MDS and both SIDs appear in Section IV-A.

IV. Experimental setup

For compatibility with our previous work [12], all experiments were executed on ALASKA II dataset with 75,000 images split into three parts (Split 1, 2, and 3), each containing 25,000 images further divided into 22k, 1k, and 2k images for training, validation, and testing. The images were developed as in [19] and then compressed with `scipy.dct` with a range of quality factors. The splits are used to study the impact of a mismatched training set for training Warden's detector and for training poolers. Alice uses the test set of Split 1 to send her secret messages in bags of size B by sampling B images without replacement. To conduct the evaluation, we utilize a specific number of test bags for each bag size. We use 2000 test bags for bag sizes 1 to 10, 1000 test bags for bag sizes $15, 20, 30$, and 250 test bags for bag sizes 60 and 120.

Our study includes JPEG qualities 98, 95, 90, 85, and 75 and two payload constraints to show how the bag gain depends on the quality factor and the way an average payload is maintained over time. Because of the sheer amount of possible combinations of the steganographer's detector, the Warden's detector, stego schemes, communication rates r, bag sizes, and spreading / pooling strategies, we limit our exposition to one content-adaptive steganographic scheme, the J-UNIWARD [18]. The average communication rate r was selected to best demonstrate the bag gain phenomenon and avoid degenerate cases when the detection is too close to random guessing or almost perfect.

A. Implementation details

For spreading, the MDS uses a single-image detector (SRNet1) trained on Split 1. Splits 2 and 3 are used by the Warden who trains her SID as another instance of SRNet (SRNet2) on Split 2. Both network detectors were pre-trained on a binary task of steganalyzing J-UNIWARD [18] (JIN pre-training [20]). The binary classifier SID was trained by drawing stego images embedded with relative payloads in terms of bpnzac selected uniformly at random from the set

$$\mathcal{P} = 0.05, 0.1, 0.2, \ldots, 0.9, 1. \tag{8}$$

The quantitative detector, which we abbreviate qSID, is also another instance of SRNet (SRNet2) pre-trained on a binary task of steganalyzing J-UNIWARD [18]. However, the qSID was trained using an L_2 loss by generating stego images in a similar manner to SID on a refined payload grid (minibatches were formed by uniformly randomly sampling images with payloads from $\mathcal{P}_{\text{fine}} \cup \{0\}$)

$$\mathcal{P}_{\text{fine}} = 0.01, 0.02, ..., 0.09, 0.1, 0.2, \ldots, 0.9, 1. \tag{9}$$

We use this form of the qSID rather than the previous state of the art [21] since the end-to-end training and architecture simplifies the computational cost of experiments. Since the design of quantitative detectors is not the subject of our current work, we refrain from including a more extensive comparison between these two quantitative detectors. We do note that when compared with the bucket estimator from [21] in the spatial domain using the same settings as described in [21], our end-to-end trained estimator achieved 10% better performance in terms of MSE and MAE (mean absolute error) scores.

The IMS was implemented by considering a given bag of B images each with N DCTs as a single large image into which the total payload was embedded using an embedding simulator. The costs were pre-computed from single images.

For MDS implementation, we employ a logistic model of response curves exactly as described in the original publication [12]. Response curves were estimated from a fine fixed grid of payloads in terms of bpnzac (9) (and thus variable from image to image in terms of bpc). We also capped the capacity of all images at 1 bpnzac because this payload is very detectable in our cover source.

Since the maximal capacity of each image is 1 bpnzac, to avoid problems with satisfying the payload constraint in bpc for small bag sizes, we ensure all images in the bag have a sufficient number of non-zero AC coefficients to meet the desired rate. Images with not enough non-zero AC coefficients are excluded from the bag formation process. Typically, this eliminates less than 1% of images from the test set of Split 1.

V. Results

In Figure 2, we show the pooler's weighted Area Under the Curve (wAUC) [19] on the test set as a function of the bag size B for two quality factors, two senders, two payload constraints, and two poolers with SID and qSID. In agreement with the analysis in [16], the bag gain is larger for the larger quality factor because higher-quality JPEGs are more likely to have flatter response curves (and thus be able to hold large payloads). The bag gain is also larger for the MDS because it assigns payloads more aggressively than IMS. Interestingly, maintaining a fixed bpc per bag leads to larger bag gain than bpnzac. This effect is explained in the appendix from the same bivariate model of response curves as in [16]. Furthermore, we note that SID outperforms qSID in all configurations. We attribute it to the fact that qSID lacks accuracy for smaller payloads. Therefore, going forward we only employ poolers equipped with the binary SID.

Figure 3 shows the bag gain as a function of JPEG quality for the IMS and MDS for both payload constraints when steganalyzed with the best pooler. The gain is generally smaller for IMS than for MDS and significantly smaller for the payload constraint in bpnzac. We explain this using a model in the appendix. Additionally, we observed that the bag size where the bag gain occurs is larger for MDS than for IMS for both forms of the payload constraint.

The bag gain does not manifest for rates that are too large (the detectability monotonically increases with B) or too small as in this case the dip "drowns" in statistical fluctuations because poolers perform as random guessers for small B. In Figure 4, we show the bag gain for quality 95, the bpc payload constraint, the MDS, and the best pooler to show the dependence on the rate r. Similar trends can be observed for other quality factors and the IMS.

It is difficult if possible to contrast the JPEG domain with the spatial domain because the embedding algorithm is different and so is the cover source. Relating to Figure 11 from [12], the bag gain for the MDS in the spatial domain with HILL as the stego scheme is significantly larger than for the IMS. In contrast, the difference between IMS and MDS is much smaller in the JPEG domain. This is because the IMS with J-UNIWARD is comparable in its aggressivity to assign payloads with the MDS (unlike the IMS with HILL). To substantiate this claim, in Figure 5, we plot the distribution of the maximal payload assigned by IMS and MDS in bag sizes corresponding to the bag gain in two different domains. The top histogram corresponds to HILL at 0.3 bpp and bag size $B = 16$ while the bottom histogram shows the distribution for J-UNIWARD at 0.1 bpc for quality factor 95 and bag size $B = 15$. The IMS with HILL in the spatial domain does not assign the largest payloads nearly as aggressively as the MDS. In contrast, the IMS and MDS with J-UNIWARD are more similar in terms of payload assignment aggressivity.

VI. Conclusions

When a batch sender maintains a fixed communication rate and assigns payloads to images based on detectability, the security is maximal when the payload is spread among a certain number of images. This gain in security is called the bag gain. In the current paper, we study this phenomenon in the JPEG domain. Our experimental findings are supported with analysis from a simple source model. In particular, the bag gain is more significant for larger JPEG qualities and for batch senders that assign the payload chunks to images in a more aggressive manner. Also, the bag gain is more pronounced when the sender maintains a fixed rate in terms of bits per DCT coefficient than per non-zero AC DCT coefficient because the latter constraint decreases the diversity of stego bags. The bag gain and the optimal bag size depend on the rate and the type of the batch sender. For the image merging sender, the gain can be up to 0.04 in terms of wAUC of the pooled detector (for high qualities) and up to 0.065 for the minimum deflection sender. Since the optimal bag size ranges from 2 to 8, the results are relevant for practitioners.

The bag gain largely disappears for very low JPEG qualities because the embedding is generally more detectable in such images and there are not enough images that are complex enough to safely hold large payloads.

Acknowledgment

The work on this paper was supported by NSF grant No. 2028119. The U.S.

Appendix

Effect of payload constraint on bag gain

As shown in [16], many experimentally observed trends in detectability as a function of bag size for the MDS can be explained with a linear model of response curves (5),

$$\varrho_i(\alpha_i) - \varrho_i(0) = b_i\alpha_i, \quad b_i \geq 0 \tag{10}$$

and a bivariate model of slopes $b_i \in \{\varepsilon, 1\}$

$$\begin{aligned} \mathbb{P}(b_i = \varepsilon) &= p \\ \mathbb{P}(b_i = 1) &= 1 - p, \end{aligned} \tag{11}$$

where $0 < \varepsilon \ll 1$ and $p \in [0, 1]$. In other words, the cover source consists of two types of images—easy to steganalyze images with slope $b = 1$ and difficult images with slope ε, which can hold a large payload with a virtually unchanged detector response. Due to the different statistical makeup of small and large bags, the detectability initially decreases, then levels off, and eventually increases due to the square root law. The advantage of the simple model (10) and (11) is a closed form[2] for MDS payloads

$$\alpha_i = \frac{r \times c(\mathbf{X})}{b_i^2 \sum_{k=1}^{B} \frac{1}{b_k^2}}, \tag{12}$$

that minimize the deflection (4). Denoting the number of images with slope ε in bag $\mathbf{X}$ with $C_\varepsilon \in \{0, 1, \ldots, B\}$,

[2]This is because we minimize a function quadratic in α_i with a linear payload constraint. Technically, the solution requires the payloads to be unbounded.

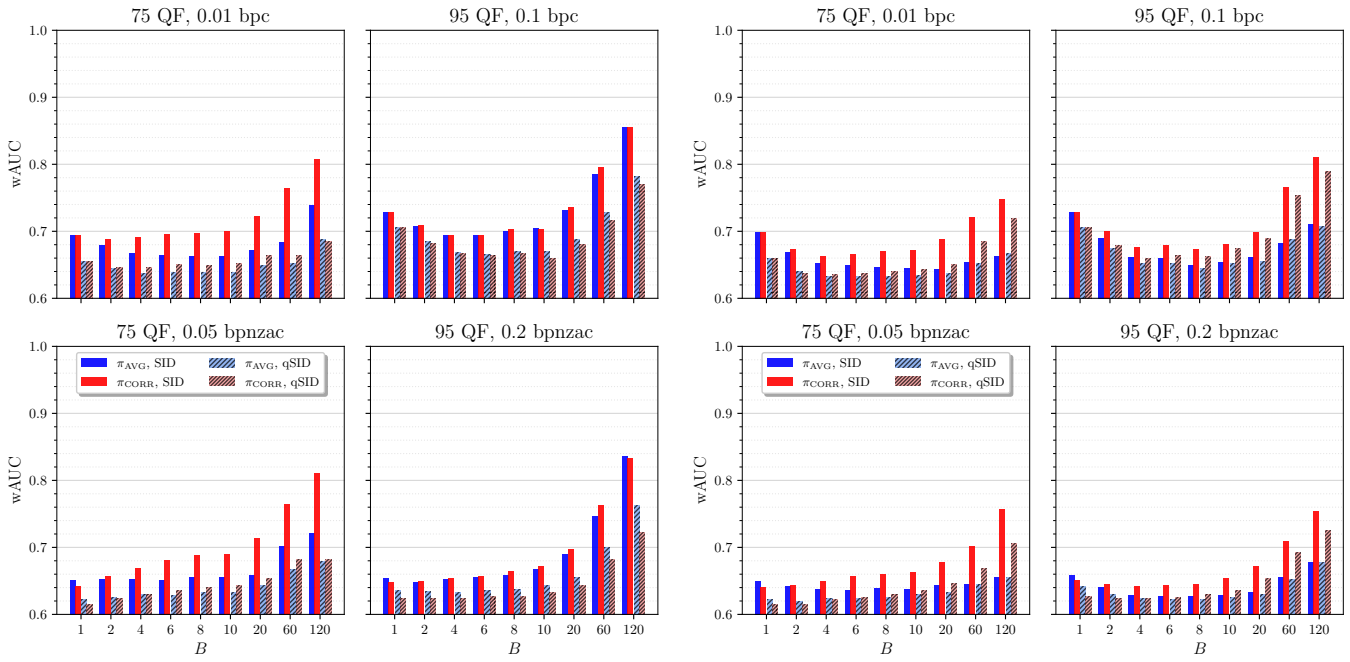

Figure 2. Pooled detector's wAUC on the test set as a function of the bag size B for IMS (left 2×2 subfigures) and MDS (right 2×2 subfigures) for quality 75 and 95 for two payload constraints, bpc (top) and bpnzac (bottom) with simple average pooler (blue, light blue) and the correlator pooler (red, light red). Solid colored bars correspond to poolers with SID, while the hatched bars to poolers with qSID.

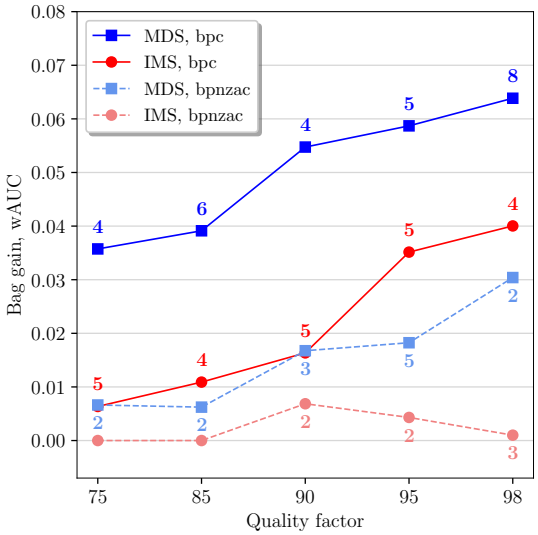

Figure 3. Bag gain in terms of wAUC as a function of JPEG quality for the IMS and MDS for both payload constraints when steganalyzed with the best pooler. Blue lines (squares) correspond to MDS, red (circles) to IMS. Each point in the plot is annotated with a bag size where the bag gain occurs. Points with no label indicate no bag gain.

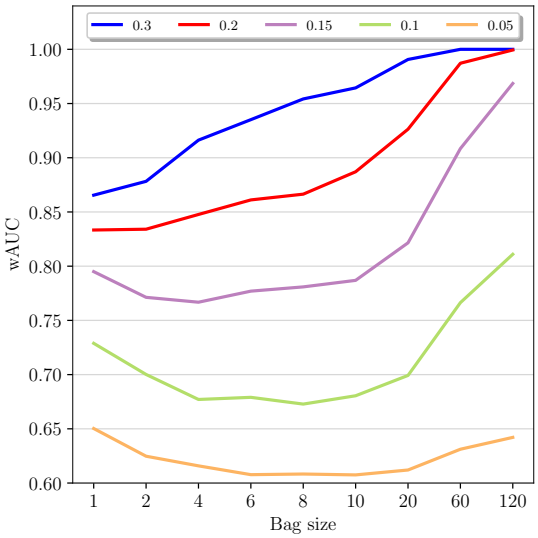

Figure 4. Pooled detector's wAUC on the test set in terms of wAUC as a function of the bag size to show the effect of the rate r on the bag gain. Payload constraint in bpc, MDS, quality 95 when steganalyzed with the best pooler.

$\sum_{k=1}^{B} \frac{1}{b_k^2} = \varepsilon^{-2} C_\varepsilon + B - C_\varepsilon$ and the deflection $\Delta^2(\mathbf{X}) = 1/\sigma^2 \sum_{i=1}^{B} b_i^2 \alpha_i^2$ becomes

$$\Delta^2(\mathbf{X}) = \frac{r^2 c^2(\mathbf{X})}{\sigma^2(\varepsilon^{-2} C_\varepsilon + B - C_\varepsilon)}. \tag{13}$$

Note that C_ε follows a binomial distribution on its support across bags $\mathbf{X}$. As shown in [16], when $c(\mathbf{X}) = B$ the trends exhibited by the expected deflection with respect to B, ε, and p qualitatively match experiments on real datasets. The bag gain will manifest as long as larger bags contain enough

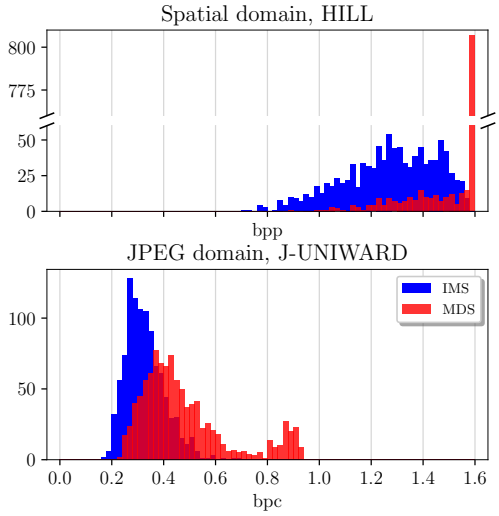

Figure 5. Distribution of the maximal payload assigned by IMS and MDS to images from bags of size $B = 16$ for spatial domain and HILL (top) and $B = 15$ for JPEG domain and J-UNIWARD (bottom). The IMS and MDS in the JPEG domain are more similar in their tendency to assign payloads aggressively than in the spatial domain.

hard-to-steganalyze images on average and the batch sender is aggressive enough in assigning the payload (IMS or MDS). The reader is referred to [16] for a detailed analysis of the above model and its ability to explain many interesting trends in detectability observed in experiments.

To understand why the bag gain is less pronounced for the payload constraint in bpnzac, note that the response curve is generally steeper (has a larger slope b) for images with fewer non-zero DCTs as such images are smoother. In contrast, highly textured / noisy images, which are more difficult to steganalyze, will have smaller slopes and a larger number of non-zero DCTs N_i. Hence, there is an approximate inverse proportionality between b_i and N_i: $N_i \propto 1/b_i^q$ for some $q > 0$. We note that for $q = 0$, $c(\mathbf{X}) \propto \sum_{i=1}^{B} 1/b_i^q = B$, which corresponds to the payload constraint in bpc (1). If we model the relationship for bpnzac with $q = 2$, $c(\mathbf{X}) = \frac{1}{N}\sum_{i=1}^{B} N_i \propto \sum_{k=1}^{B} \frac{1}{b_k^2}$, and (13) simplifies

$$\Delta^2(\mathbf{X}) \propto \sum_{i=1}^{B} b_i^2 \frac{c^2(\mathbf{X})}{b_i^4 \left(\sum_{k=1}^{B} \frac{1}{b_k^2}\right)^2} \tag{14}$$

$$\propto \sum_{i=1}^{B} \frac{1}{b_i^2} = \varepsilon^{-2} C_\varepsilon + (B - C_\varepsilon). \tag{15}$$

Since C_ε has a binomial distribution, the expectation of the deflection increases *monotonically* w.r.t. bag size B, suppressing the manifestation of the bag gain.

References

[1] A. D. Ker, "Batch steganography and pooled steganalysis," in *Information Hiding, 8th International Workshop* (J. L. Camenisch, C. S. Collberg, N. F. Johnson, and P. Sallee, eds.), vol. 4437 of Lecture Notes in Computer Science, (Alexandria, VA), pp. 265–281, Springer-Verlag, New York, July 10–12, 2006.

[2] A. D. Ker, "Batch steganography and the threshold game," in *Proceedings SPIE, Electronic Imaging, Security, Steganography, and Watermarking of Multimedia Contents IX* (E. J. Delp and P. W. Wong, eds.), vol. 6505, (San Jose, CA), pp. 04 1–13, January 29–February 1, 2007.

[3] A. D. Ker, "Perturbation hiding and the batch steganography problem," in *Information Hiding, 10th International Workshop* (K. Solanki, K. Sullivan, and U. Madhow, eds.), vol. 5284 of Lecture Notes in Computer Science, (Santa Barbara, CA), pp. 45–59, Springer-Verlag, New York, June 19–21, 2008.

[4] A. D. Ker and T. Pevný, "Batch steganography in the real world," in *Proceedings of the 14th ACM Multimedia & Security Workshop* (J. Dittmann, S. Craver, and S. Katzenbeisser, eds.), (Coventry, UK), pp. 1–10, September 6–7, 2012.

[5] X. Hu, J. Ni, W. Zhang, and J. Huang, "Efficient JPEG batch steganography using intrinsic energy of image contents," *IEEE Transactions on Information Forensics and Security*, vol. 16, pp. 4544–4558, 2021.

[6] L. Li, W. Zhang, C. Qin, K. Chen, W. Zhou, and N. Yu, "Adversarial batch image steganography against CNN-based pooled steganalysis," *Signal Processing*, vol. 181, pp. 107920–107920, 2021.

[7] V. Sedighi, R. Cogranne, and J. Fridrich, "Practical strategies for content-adaptive batch steganography and pooled steganalysis," in *Proceedings IEEE, International Conference on Acoustics, Speech, and Signal Processing*, March 5–9, 2017.

[8] M. Sharifzadeh, M. Aloraini, and D. Schonfeld, "Adaptive batch size image merging steganography and quantized Gaussian image steganography," *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 867–879, 2020.

[9] T. Pevný and I. Nikolaev, "Optimizing pooling function for pooled steganalysis," in *IEEE International Workshop on Information Forensics and Security*, (Rome, Italy), pp. 1–6, November 16–19, 2015.

[10] N. Zhong, Z. Qian, Z. Wang, X. Zhang, and X. Li, "Batch steganography via generative network," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 31, pp. 88–97, January 2021.

[11] A. Zakaria, M. Chaumont, and G. Subsol, "Pooled steganalysis in JPEG: how to deal with the spreading strategy?," in *IEEE International Workshop on Information Forensics and Security (WIFS)*, pp. 1–6, 2019.

[12] Y. Yousfi, E. Dworetzky, and J. Fridrich, "Detector-informed batch steganography and pooled steganalysis," in *The 10th ACM Workshop on Information Hiding and Multimedia Security* (J. Butora, B. Tondi, and C. Veilhauer, eds.), (Santa Barbara, CA), ACM Press, 2022.

[13] M. Aloraini, M. Sharifzadeh, and D. Schonfeld, "Quantized Gaussian JPEG steganography and pool steganalysis," *IEEE Access*, vol. 10, pp. 38031–38044, 2022.

[14] X. Liao and J. Yin, "Two embedding strategies for payload distribution in multiple images steganography," in *IEEE ICASSP*, pp. 1982–1986, April 15-20 2018.

[15] A. D. Ker, "The square root law of steganography," in *The 5th ACM Workshop on Information Hiding and Multimedia Security* (M. Stamm, M. Kirchner, and S. Voloshynovskiy, eds.), (Philadelphia, PA), ACM Press, June 20–22, 2017.

[16] E. Dworetzky and J. Fridrich, "Explaining the bag gain in batch steganography," *IEEE Transactions on Information Forensics and Security*, vol. 18, pp. 3031–3043, 2023.

[17] T. H. Thai, R. Cogranne, and F. Retraint, "Camera model identification based on the heteroscedastic noise model," *Image Processing, IEEE Transactions on*, vol. 23, pp. 250–263, Jan 2014.

[18] V. Holub, J. Fridrich, and T. Denemark, "Universal distortion design for steganography in an arbitrary domain," *EURASIP Journal on Information Security, Special Issue on Revised Selected Papers of the 1st ACM IH and MMS Workshop*, vol. 2014:1, 2014.

[19] R. Cogranne, Q. Giboulot, and P. Bas, "ALASKA–2: Challenging academic research on steganalysis with realistic images," in *IEEE International Workshop on Information Forensics and Security*, (New York, NY), December 6–11, 2020.

[20] J. Butora, Y. Yousfi, and J. Fridrich, "How to pretrain for steganalysis," in *The 9th ACM Workshop on Information Hiding and Multimedia Security* (D. Borghys and P. Bas, eds.), (Brussels, Belgium), ACM Press, 2021.

[21] M. Chen, M. Boroumand, and J. Fridrich, "Deep learning regressors for quantitative steganalysis," in *Proceedings IS&T, Electronic Imaging, Media Watermarking, Security, and Forensics 2018* (A. Alattar and N. D. Memon, eds.), (San Francisco, CA), January 29–February 1, 2018.